Why is Insurance Good? An Example

Jon Bakija, Williams College (Revised October 2013)

Introduction

The United States government is, to a rough approximation, "an insurance company with an army."[1] That is because the majority of federal government spending represents some form of "social insurance." Rising health care costs and an aging population are expected to make the share of government spending devoted to social insurance in the U.S. grow rapidly in the future, and social insurance is already a much larger share of government in essentially every other rich country than it is in the U.S. So to understand modern political economy, it is absolutely critical to have a solid understanding of why insurance is good. That insurance has value is a non-obvious insight, because on average, people lose money when they buy insurance, due to the need to cover the costs of the insurance company. So why would people ever voluntarily buy insurance? How could it possibly make them better off? Let's consider a simple example that gets the idea across.

Set-Up of the Example

Suppose you face a risky situation. To keep the math simple, suppose there is a 50% probability that you will have good luck, in which case your consumption will be $90,000, and a 50% probability that you will have bad luck, in which case your consumption will be $10,000. You can think of the difference as representing some catastrophic expense that occurs randomly, such as a large health expenditure, in which case by "consumption" we mean consumption aside from the catastrophic expense. Or you can think of it as a drop in consumption arising, for example, from an unexpected job loss. Further suppose that your utility equals the square root of your consumption, which is a mathematically convenient function that exhibits diminishing marginal utility of consumption. So to summarize:

Utility = $U(C) = \sqrt{C} = C^{1/2}$. Note that "$U(C)$" means "utility as a function of consumption".
Consumption if good luck, Cg = $90,000
Consumption if bad luck, Cb = $10,000
Probability of good luck, Pg = 50%
Probability of bad luck, Pb = 50%

The "expected value" of something is just its value on average. You can calculate the expected value of something by multiplying each possible value of the variable by its probability, and then summing up. In general, the symbol for "expected value of something" is E(something).

[1] The original source of this quote is Peter Fisher, from when he was Undersecretary of the Treasury, in 2002 <http://economistsview.typepad.com/economistsview/2013/01/who-first-said-the-us-is-an-insurance-company-with-an-army.html>.

The average, or "expected," value of consumption, $E(C)$, when facing this risky situation, is:

$E(C) = Pb \times Cb + Pg \times Cg = 0.5 \times \$10{,}000 + 0.5 \times \$90{,}000 = \$50{,}000$.

We'll denote this as $E(C)_{NO\ INSURANCE}$, to remind us that this is your expected value of consumption when you do not have insurance.

The average, or "expected," value of utility, E(U), when facing this risky situation, is:

$E(U) = Pb \times U(Cb) + Pg \times U(Cg) = 0.5 \times \sqrt{10{,}000} + 0.5 \times \sqrt{90{,}000} = 200$ utils.

We'll denote this as $E(U)_{NO\ INSURANCE}$, to remind us that this is your expected value of utility when you do not have insurance.

A rational person should behave so as to maximize his or her own expected value of *utility*, not his or her own expected value *consumption*. Utility is your level of happiness. That's what you care about. Consumption is only instrumental to utility (happiness). So you don't care directly about the dollar value of consumption, but instead you care about the amount of utility provided by that dollar value of consumption, and you should behave so as to maximize expected utility.

Why Risk is Bad

Diminishing marginal utility implies that on average, your expected level of happiness (utility) would be higher if you could somehow ensure that you would always get the expected value of your consumption with certainty, instead of having consumption that was much lower than its expected value in certain "bad luck" states of the world that might occur, and much higher than its expected value in other "good luck" states of the world that might occur. This is true because, if there is diminishing marginal utility from consumption, the gain in utility from increasing consumption above its expected value in the good luck state of the world is smaller than the loss of utility from reducing consumption below its expected value by the same dollar amount in the bad luck state of the world. Figure 1 below depicts the scenario described in the example above, and demonstrates the conclusion that "risk is bad" in a more systematic way.

If you could somehow maintain the expected value of your consumption of $50,000 *with certainty*, then your utility would be $U(50{,}000) = \sqrt{50{,}000} = 224$ utils. That is better than the 200 utils of utility that you expect to get, on average, from the risky situation where half the time you consume $10,000 (yielding a utility of 100) and the other half of the time you consume $90,000 (yielding a utility of 300). Figure 1 illustrates this as follows. Compared to situation where you get the expected value of your consumption ($50,000) with certainty, which yields a utility of 224, the gain in utility from an increase in consumption from $50,000 to $90,000 is only (300-224) = 76 utils, whereas the loss in utility from an equal-dollar decrease in consumption, from $50,000 to $10,000, is a much larger (224 – 100) = 124 utils. Going up by $40,000 relative to a given starting point increases your happiness by much less than going down by $40,000 from the same starting point reduces your happiness. That will be true as long as your utility function exhibits diminishing marginal utility of consumption. A person whose utility function exhibits diminishing marginal utility of consumption will thus dislike risk, will be willing to pay to avoid it, and will demand compensation to accept it – we call such a person "risk averse."

Figure 1 -- Insurance is good: example with 50% chance of bad luck and $U=C^{1/2}$

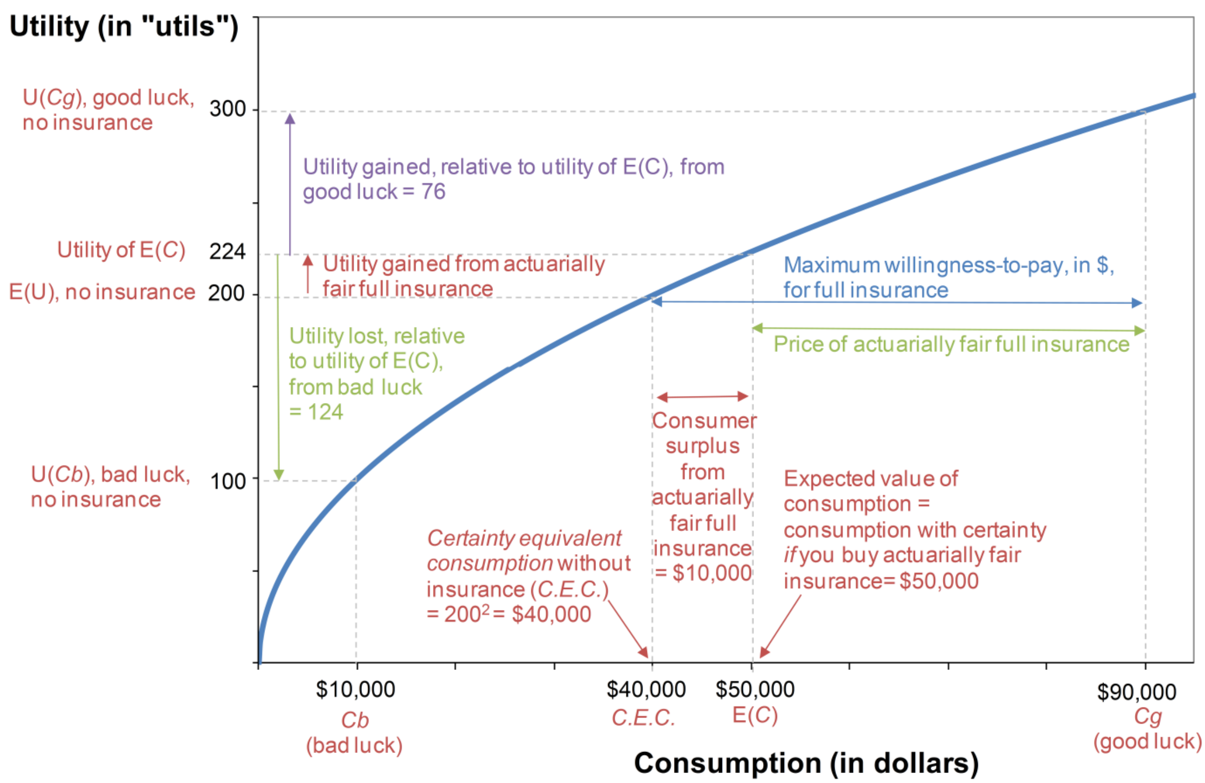

Note that the argument for why risk is bad relies on the same principle that is behind the utilitarian argument for why redistribution from rich to poor can increase the sum of utilities in society. In our stylized example, you prefer having $50,000 for sure to having $10,000 with 50 percent probability and $90,000 with 50 percent probability. If you could somehow redistribute $40,000 away from yourself in the possible state of the world where you have $90,000, and send it to yourself in the possible state of the world where you only have $10,000, the probability-weighted sum of your utilities across the two states of the world (that is, your expected utility) would go up. That is more or less what insurance does. The difference from redistribution is that you are willing to enter into an insurance contract voluntarily, because you don't know in advance whether the good luck or bad luck state of the world will apply to you. As a result, under certain conditions, private insurance companies can make money supplying insurance.

How Insurance Works

In the absence of administrative costs, an insurance company that has accurate information on your probability of bad luck can sell you an insurance policy that guarantees that you will get the *expected value* of your consumption with certainty. (We'll discuss how to deal with administrative costs later). Imagine there are a large number of people exactly like the "you" described in the example above. So, each of them has a probability of bad luck of 50 percent. Further suppose that their probabilities of bad luck are *independent*, meaning that if one person turns out to have bad luck, that has no effect one way or the other on anyone else's likelihood of having bad luck -- an independent coin flip determines whether each individual turns out to have good luck or bad luck. Then the insurance company can offer the following insurance contract. Before you know whether you will have good luck or bad luck, you are offered the opportunity to pay a price (the price is also known as a "premium") for an insurance policy. If it turns out that you have bad luck, the insurance company pays you a "benefit" of $80,000. If it turns out that you have good luck, the insurance company does not pay you any benefit at all. This is an example of "full" insurance, because the insurance benefit payment is exactly equal to the difference between consumption in the good luck state of the world ($90,000) and consumption in the bad luck state of the world ($10,000) in the absence of insurance. Full insurance completely eliminates the customer's risk – the customer's consumption will be identical, and equal to $90,000 minus the price paid for insurance, regardless of luck. The insurance company also offers a similar insurance contract to everyone else that is like you.

In that case, what is the expected value of the marginal cost to the insurance company of supplying full insurance to one more customer? In the absence of administrative costs, the expected marginal cost is the probability of bad luck times the benefit payment, or 0.5×$80,000 = $40,000. In a competitive market, the price of insurance would be driven down to that marginal cost, so full insurance would sell for a price of $40,000, and each insurance company would just break even, on average. When the price of full insurance is equal to the probability of the bad outcome, times the benefit payout in the event of the bad outcome, we call that "actuarially fair full insurance." When insurance companies without administrative costs sell actuarially fair full insurance to a large enough pool of customers with independent probabilities of bad luck, then the "law of large numbers" suggests that the probability that each insurance company will break even is very high. Intuitively, if you flip a coin a very large number of times, the probability that the proportion of flips that are "heads" will be very close to 50% is very close to one, and that probability gets closer and closer to one the more times you flip the coin (which corresponds to having a larger pool of customers).[2]

[2] If the probabilities were not independent, the math would get harder, but insurance could still work as long as the outcomes are not *perfectly* correlated – that is, not *everyone* has bad luck at the same time – or if the insurance company can pool risks across time (for example by saving), or can pool the risks across a larger number of customers who do not have perfectly correlated risks. In general, correlated risks would raise the break-even price of insurance, unless the insurance company can somehow insure itself against the correlated portion of the risk. That is what the "reinsurance" business is about.

If you buy the actuarially fair full insurance, in the good luck state of the world your consumption is $90,000 minus the $40,000 premium = $50,000, and in the bad luck state of the world your consumption is $10,000 plus the $80,000 insurance benefit minus the $40,000 premium = $50,000. Thus, if you buy the insurance, you are guaranteed to consume $50,000 for sure, regardless of how your luck turns out. Recall that $50,000 was the expected value of your consumption in the uncertain situation. Thus, actuarially fair full insurance enables you to maintain the expected value of your consumption with certainty, and eliminates risk.

The Utility Gain from Insurance

How much better off does the insurance make you? As noted above, the utility from getting $50,000 with certainty is approximately 224 utils. That is your level of utility if you buy the insurance. If you do not buy the insurance, your expected utility from a 50% chance of $10,000 of consumption and a 50% chance of $90,000 of consumption is only 200 utils. So buying the insurance makes you better off by 24 utils. Of course, if you turn out to have good luck, you would have been better off not buying the insurance, but there was no way for you to know that in advance. Protecting yourself from risk by buying actuarially fair full insurance makes you better off *on average*, not in each particular possible state of the world.

The Economic Surplus Gain from Insurance

As we will see, insurance markets often fail due to asymmetric information. That is a market failure, and if there is a market failure, the outcome is economically inefficient. As you know, economic efficiency is about maximizing "economic surplus," or *dollar-valued* net benefits. So to think about the economic efficiency consequences of insurance market failure, we need to be able to measure how much better off the insurance makes people *in dollars*. The principles are the same as for any other good. So in the absence of market failure, the economic surplus from insurance equals the sum of producer surplus and consumer surplus. We also need to figure out the benefits and costs measured in dollars in order to understand under what conditions a person will or will not voluntarily buy insurance, and to work out supply and demand in the market, because consumers pay for things with dollars, not with utils.

The producer surplus from each insurance contract equals the price the insurance company receives for selling it, minus the marginal cost of supplying it. In our example, the marginal cost of supplying insurance to a customer with a particular probability of bad luck is the same, no matter how many of these customers are served. So the marginal cost is constant. The height of the supply curve is equal to the marginal cost, so the supply curve will be horizontal, and producer surplus will be zero in the competitive market equilibrium.

The consumer surplus from each insurance contract equals the consumer's maximum willingness-to-pay for that insurance, minus its price. Maximum willingness-to-pay always determines the height of the demand curve. So to figure out consumer surplus from insurance, we need to figure out maximum willingness-to-pay for the insurance.

The key insight that enables us to determine the maximum willingness-to-pay for insurance is that the expected value of utility associated with a risky situation can be translated into a *dollar-valued* measure of expected well-being. To do that, we have to answer the question "what level of consumption, if you knew you would have it with *certainty*, would give you the same utility as the *expected* value of your utility in the risky situation?" We call the answer to that question the <u>certainty equivalent consumption</u> (or "*C.E.C.*" for short). By definition, the utility derived from "certainty equivalent consumption" is equal the expected utility from the risky situation (the situation without insurance). Recall that in our example, utility = $\sqrt{C}$, and that $E(U)_{\text{NO INSURANCE}}$ = 200 utils. Therefore:

$\sqrt{C.E.C}$ = 200 utils

We can solve for *C.E.C.* by squaring both sides of the equation above. This gives us *C.E.C.* = $40,000. In words, the risky situation (the situation without insurance) provides the same level of utility, on average, as if the person had 200^2 = $40,000 of consumption for certain. $40,000 is the certainty equivalent consumption.

What is the maximum amount you would be willing to pay for full insurance? In this example, if you buy insurance, your consumption will always be $90,000 minus the price of the insurance. If you have good luck, you get the consumption you would have had anyway ($90,000), but you lose the price you paid for the insurance, and get no benefit payment from the insurance company. If you have bad luck, you get the consumption you would have had with bad luck ($10,000), plus the insurance benefit payment of $80,000, minus the price you paid for the insurance. So if you buy the insurance, your consumption is guaranteed to be $90,000 minus the price you pay for the insurance, regardless of how your luck turns out. If you do *not* buy the insurance, your expected level of well-being, measured in dollars, is your certainty equivalent consumption.

So, the buying the insurance will make you better off if and only if:

(Guaranteed consumption with the insurance) >
(Certainty equivalent consumption without the insurance)

Since the guaranteed consumption with the insurance in this example is $90,000 minus the price, and since the certainty equivalent consumption is $40,000, the insurance only makes you better off if and only if:

$90,000 – price > $40,000

The price that just makes you indifferent between buying insurance and not buying insurance is the one that would make your guaranteed consumption with the insurance just equal to your certainty equivalent consumption without the insurance. That price is your maximum willingness to pay for the insurance. If the price was any higher, buying the insurance would make you worse off. Therefore, if MWTP stands for maximum willingness to pay, then:

$\$90{,}000 - MWTP = \$40{,}000$

Solving the equation above for *MWTP* leads to the conclusion that your maximum willingness to pay for the insurance is $50,000. If the premium goes any higher than $50,000, you will not buy it.

What is your consumer surplus from the insurance? Just as with any other good, it is maximum willingness-to-pay minus price. As we showed above, if this insurance is sold in a competitive market and there are no administrative costs, the price of the insurance is $40,000. So your consumer surplus from the insurance is $50,000 - $40,000 = $10,000.

Since producer surplus is $0, total economic surplus from the insurance is just the consumer surplus of $10,000. "Deadweight loss" is defined as the loss of economic surplus that occurs because of an economically inefficient situation. If asymmetric information causes the insurance market to fail, and as a result you cannot buy insurance at a price that is mutually beneficial to you and the insurance company, then the deadweight loss from that missing insurance is $10,000, which is equal to the economic surplus we *could have had* if the market did not fail.

Factoring in Administrative Costs

In the real world, insurance companies face additional costs above and beyond the expected costs of paying out benefits. These additional costs include, for example, the wages of actuaries and of workers involved in administering the insurance contracts, the costs of computers and rent on office buildings, and the need to cover the opportunity costs the owners incur by putting their savings into the firm (i.e., they could alternatively have earned interest on those savings). We'll call those costs "administrative costs." Administrative costs increase the marginal cost of supplying the insurance, so the price (premium) has to go up to cover the additional costs. If we change our example to include an administrative cost of $1,000 per customer, the marginal cost of supplying the insurance rises to $41,000, and in a competitive market, so does the price. Note that the administrative cost in this example is incurred for every customer, not just the ones who have bad luck. The producer surplus in a competitive market is still zero because we still have constant marginal costs (i.e the marginal cost does not depend on the number of customers), and the price equals marginal cost. Those marginal costs include the normal profit or opportunity cost of capital, which we've included in "administrative costs" here. So insurance companies in a perfectly competitive market with constant marginal costs still earn "accounting profits" (which include profits that cover the normal opportunity cost of capital), and are perfectly happy to supply the insurance as a result, but they do not earn any "economic profit" (economic profits are profits above and beyond the amount needed to cover the ordinary opportunity cost of capital and other costs). That's what we mean when we say producer surplus is zero. The consumer surplus is still *MWTP* – price, which now equals

$50,000 - $41,000 = $9,000. As long as the administrative cost per customer is lower than the difference between *MWTP* and the marginal cost of actuarially fair full insurance, there is positive economic surplus to be had from the insurance contract, and room for a deal to be made between consumers and producers that makes both better off.

Summary and Conclusion

To summarize, the following steps are involved in calculating economic surplus from insurance, and in figuring out who will buy insurance when it is offered at a particular price.

1. Calculate a person's <u>certainty equivalent consumption</u> in the absence of insurance.

2. Calculate the person's <u>maximum willingness-to-pay</u> for the insurance. This is the price of insurance that would make the guaranteed consumption with the insurance exactly equal to the certainty equivalent consumption without the insurance. Remember that to find "guaranteed consumption with the insurance," you need to subtract off the price paid for the insurance from the level of consumption the insurance would guarantee the person if he or she got the insurance for free.

3. Calculate the <u>marginal cost</u> and the <u>price</u> of the insurance. The marginal cost equals the administrative cost per customer, plus the expected value of the benefit payout per customer. In a competitive market, the insurance company sets the price of the insurance equal to marginal cost.

    a. If there are multiple risk types, and the insurance company can tell them apart, then the insurance will be sold at a different price to each customer, with price equal to marginal cost for that type of customer.

    b. If there are multiple risk types and the insurance company *cannot* tell them apart, then each insurance company will charge a price equal to the average marginal cost of all the customers it expects to attract. If at first it gets this wrong, it will eventually adjust its premium to reflect the marginal cost of the customers it actually attracts, so that it can break even.

4. An individual will voluntarily buy insurance if and only if maximum willingness-to-pay is greater than or equal to the price at which he or she is offered insurance.

5. Consumer surplus for each customer, if he or she buys the insurance, is his or her maximum willingness-to-pay minus price.

6. Producer surplus from each customer, if he or she buys the insurance, is price minus marginal cost.

7. Economic surplus associated with each customer is the sum of consumer and producer surplus.

8. Deadweight loss is the amount of economic surplus lost if a market failure (such as asymmetric information) prevents a mutually beneficial insurance contract from being purchased.

The basic insights are easier to see when the probability of bad luck is 50%, but do not depend on it. The key point is that as long as administrative costs are low enough and firms can charge a premium to each customer that reflects the customer's true level of risk, then the insurance will be mutually beneficial to customer and insurance company, and will increase both utility and economic surplus in society relative to the absence of insurance. Diminishing marginal utility is what makes this work. When there is diminishing marginal utility, moving consumption away from yourself in high-consumption states of the world, where marginal utility is low, and towards yourself in low-consumption states of the world, where marginal utility is high, can increase the expected value of your utility. So you are voluntarily willing to do things that make that happen.